



Oracle

1Z0-058 Exam

Oracle Real Application Clusters 11g Release 2 and Grid Infrastructure Administration

Thank you for Downloading 1Z0-058 exam PDF Demo

You can Buy Latest 1Z0-058 Full Version Download

<https://www.certkillers.net/Exam/1Z0-058>

<https://www.certkillers.net>

Question: 1

Which three actions would be helpful in determining the cause of a node reboot?

- A. determining the time of the node reboot by using the uptime command and subtracting the up time from the current system time
- B. looking for messages such as "Oracle CSSD failure. Rebooting for cluster integrity" in /var/log/messages
- C. using the crsctl command to view tracing information
- D. inspecting the ocssd log for "Begin Dump" or "End Dump" messages
- E. inspecting the database alert log for reboot messages

Answer: A, B, D

Explanation:

Determining Which Process Caused Reboot

Log File Locations for Processes Causing Reboots.

- oclskd
 - <Grid_Home>/log/<hostname>/client/oclskd.log
- ocssd
 - /var/log/messages
 - <Grid_Home>/log/<hostname>/cssd/ocssd.log
- csssdagent
 - <Grid_Home>/log/<hostname>/agent/ohasd/oracss_dagent_root
- cssdmonitor
 - <Grid_Home>/log/<hostname>/agent/ohasd/oracss_dmonitor_root
- hangcheck-timer
 - /var/log/messages

First, determine the time of the node reboot by using the uptime command and subtracting the up time from the current system time. The reboot time will be used when examining log files. When the OCSSD daemon is responsible for rebooting a node, a message similar to "Oracle CSSD failure.

Rebooting for cluster integrity" is written into the system messages log at /var/log/messages.

The cssd daemon log file that is located at <Grid_Home>/log/<hostname>/cssd/ocssd.log may also contain messages similar to "Begin Dump" or "End Dump" just before the reboot.

If hangcheck-timer is being used, it will provide message logging to the system messages log when a node restart is initiated by the module. To verify whether this process was responsible for the node reboot, examine the /var/log/messages file and look for an error message similar to: "Hangcheck: hangcheck is restarting the machine."

Other useful log files include the Clusterware alert log in <Grid_home>/log/<hostname> and the lastgasp log in /etc/oracle/lastgasp or /var/opt/oracle/lastgasp.

If no indication of which process caused the reboot can be determined from these files, additional debugging and tracing may need to be enabled.

Note: The oclsomon and the oprocd background processes have been eliminated in Oracle

Database 11g Release 2.
D60488GC11
Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 6 - 21

Question: 2

After Oracle Grid Infrastructure has been installed, you should take a few moments to verify the installation. Which two actions would be useful in verifying the installation?

- A. Run the crsctl status resource –t command to confirm that all necessary cluster resources are online.
- B. Use the operating system utilities to verify that your SCAN addresses are being properly resolved.
- C. Start Oracle Enterprise Manager and check all monitored targets.
- D. Run the cluvfy comp nodecon –n all –verbose command to verify the entire Grid Infrastructure installation.

Answer: A, B

Explanation:

Verifying the Grid Infrastructure Installation

Execute the crsctl command as shown in the slide to confirm that all cluster resources are up and running.

```
root@racnode01 ~]# /u01/app/11.2.0/grid/bin/crsctl stat res -t
```

In addition, you should confirm that your DNS is properly forwarding address requests for your application and SCAN VIPs to your GNS and that they are resolved properly. You can do this with dig.

Execute the dig command with DNS and VIP addresses as shown:

```
# dig @myDNS.example.com racnode01-cluster01.example.com
...
;; QUESTION SECTION:
;racnode01-vip.cluster01.example.com. IN A
;; ANSWER SECTION:
racnode01-vip.cluster01.example.com. 120 IN A
192.0.2.103
...
# dig @myDNS.example.com cluster01-scan.cluster01.example.com
...
;; QUESTION SECTION:
;cluster01-scan.cluster01.example.com. IN A
;; ANSWER SECTION:
cluster01-scan.cluster01.example.com. 120 IN A 192.0.2.248
cluster01-scan.cluster01.example.com. 120 IN A 192.0.2.253
cluster01-scan.cluster01.example.com. 120 IN A 192.0.2.254
```

D60488GC11
Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 2 – 44

Question: 3

Which two network addresses are required to be static, non-dhcp addresses when using the Grid Naming?

- A. GNS VIP Address
- B. SCAN VIP Address
- C. Node VIP Address
- D. Node Public Address
- E. Node Private Address

Answer: A, D

Explanation:

2.6.2 IP Address Requirements

Before starting the installation, you must have at least two network adapters configured on each node: One for the private IP address and one for the public IP address.

You can configure IP addresses with one of the following options:

Dynamic IP address assignment using Oracle Grid Naming Service (GNS). If you select this option, then network administrators assign static IP address for the physical host name and dynamically allocated IPs for the Oracle Clusterware managed VIP addresses.

Oracle® Grid Infrastructure Installation Guide

Implementing GNS

To implement GNS, you must collaborate with your network administrator to obtain an IP address on the public network for the GNS VIP. DNS uses the GNS VIP to forward requests for access to the cluster to GNS. You must also collaborate with your DNS administrator to delegate a domain to the cluster. This can be a separate domain or a subdomain of an existing domain. The DNS server must be configured to forward all requests for this new domain to the GNS VIP. Since each cluster has its own GNS, it must be allocated a unique domain of which to be in control.

Oracle® Clusterware Administration and Deployment Guide

Question: 4

You are in the planning stages for upgrading your Oracle RAC database from Oracle Database 10g Release 2 to Oracle Database 11g Release 2 to run under the Oracle Grid Infrastructure. You decide to use an administrator-managed configuration because the cluster is fairly small. Which statement is correct about this configuration?

- A. A parent pool of the GENERIC server pool will be used.
- B. You must define a new server pool called MANUAL.
- C. A subpool of the GENERIC server pool will be used.
- D. A subpool of the FREE server pool will be used.

Answer: C

Explanation:

With Oracle RAC 11g Release 2, there is only a database resource define to Oracle Clusterware. This resource will contain the Oracle Home, the spfile, the server pool(s) and the ASM diskgroup(s) required for the database. The database resource will have a weak start dependency on the VIP

which means it will try to start the VIP for the node when the instance starts however if the VIP does not start successfully, the instance will still be started. When reviewing the database resource for an administrator Managed database, you will see a server pool defined with the same name as the Oracle database. This pool will be part of a special Oracle defined server pool called Generic. The Generic server pool is managed by Oracle to support Administrator Managed databases. The server pools that are part of Generic will be automatically created or removed when you add or remove an Administrator Managed database.

Oracle Database 11gR2 Upgrade Companion (Version 2.70)

Question: 5

The Instance Initialization parameters are set to:

DB_CREATE_FILE_DEST = +DATA
DB_CREATE_ONLINE_LOG_DEST_1 = +LOGS
DB_CREATE_ONLINE_LOG_DEST_2 = + FRA

The SQL* Plus command "ALTER DATABASE ADD LOGFILE;" will create:

- A. a new log file in the +DATA disk group, or a log file in the + FRA disk group, if +DATA is not available
- B. a new log file in the +DATA disk group and a log file in the + FRA disk group
- C. a new log file in the +LOGS disk group and a log file in the + FRA disk group
- D. a new log file in the +LOGS disk group, or a log file in the +FRA dls* available
- E. a new log file in the +DATA disk group, a log file in the +LOGS disk group, and a log file in the +FRA disk group
- F. a new log file in the +LOGS disk group, or a log file in the +FRA disk group, if +LOGS is not available

Answer: C

Explanation:

The following table lists the initialization parameters that enable the use of Oracle Managed Files.

Initialization Parameter

DB_CREATE_FILE_DEST

Defines the location of the default file system directory or Oracle ASM disk group where the database creates datafiles or tempfiles when no file specification is given in the create operation. Also used as the default location for redo log and control files if DB_CREATE_ONLINE_LOG_DEST_n are not specified.

DB_CREATE_ONLINE_LOG_DEST_n

Defines the location of the default file system directory or Oracle ASM disk group for redo log files and control file creation when no file specification is given in the create operation. By changing n, you can use this initialization parameter multiple times, where n specifies a multiplexed copy of the redo log or control file. You can specify up to five multiplexed copies.

DB_RECOVERY_FILE_DEST

Defines the location of the Fast Recovery Area, which is the default file system directory or Oracle ASM disk group where the database creates RMAN backups when no format option is used, archived logs when no other local destination is configured, and flashback logs. Also used as the default location for redo log and control files or multiplexed copies of redo log and control files if DB_CREATE_ONLINE_LOG_DEST_n are not specified.

Oracle® Database Administrator's Guide

Question: 6

Which three statements define a cluster?

- A. is a group of independent, but interconnected computers that act as a single system
- B. can be deployed to increase availability and performance
- C. can be deployed to balance a dynamically changing workload
- D. should appear to an application as multiple servers

Answer: A, B, C

Explanation:

Oracle Real Application Clusters

- Oracle Real Application Clusters (Oracle RAC) is a database clustering technology whose shared storage capabilities allow multiple machines to work in parallel on the same data, reducing processing time significantly. Oracle RAC also offers resilience, allowing processing to continue in the event of one or more machines being unavailable because of planned or unplanned downtime.

Computer cluster

- A computer cluster consists of a set of loosely connected computers that work together so that in many respects they can be viewed as a single system.
- The components of a cluster are usually connected to each other through fast local area networks, each node running its own instance on an operating system. Computer clusters emerged as a result of convergence of a number of computing trends including the availability of low cost microprocessors, high speed networks, and software for high performance distributed computing.
- Clusters are usually deployed to improve performance and availability over that of a single computer, while typically being much more cost-effective than single computers of comparable speed or availability

Question: 7

You want to create an ACFS on an ADVM volume using a shell script and the appropriate command-line utilities. These are the requirements:

1. The dynamic volume file must use space in the VOLFILE disk group with a size of 500 M and be called prodvol.
2. The mount point called /acfs already exists.

Which four steps must be performed to achieve this?

- A. As the Grid Infrastructure owner, run `mount -t acfs /dev/asm/prodvol-417 /acfs` to mount the file system.
- B. As the Grid Infrastructure owner, run `asmcmd volinfo -d volfile prodvol` to determine the volume information.
- C. As the Grid Infrastructure owner, run `asmcmd volcreate -d volfile -s 500M prodvol` to create the volume file.
- D. As the Grid Infrastructure owner, run `mkfs -t acfs /dev/asm/prodvol-417` to create the file system.
- E. As root, run `mount -t acfs /dev/asm/prodvol-417 /acfs` to mount the file system.
- F. As root, run `mkfs -t acfs /dev/asm/prodvol-417` to create the file system.

Answer: B, C, D, E

Explanation:

Explanation:

Reference:

Creating an ACFS Volume

Create the volume:

```
$ asmcmd volcreate -G DATA -s 100M testvol
```

View the volume information:

```
$ asmcmd volinfo -G DATA testvol
```

Make a mount point directory:

```
$ mkdir /u01/app/oracle/acfsdata/testvol
```

Make the file system (as root):

```
# mkfs -t acfs /dev/asm/testvol-403
```

Mount the file system to the mount point:

```
# mount -t acfs /dev/asm/testvol-403 \
/u01/app/oracle/acfsdata/testvol
```

D60488GC11

Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 10 - 13

Create a file system with the Oracle ACFS mkfs command.

Create a file system using an existing volume device.

For example:

```
$ /sbin/mkfs -t acfs /dev/asm/volume1-123
```

```
mkfs.acfs: version = 11.2.0.1.0.0
```

```
mkfs.acfs: on-disk version = 39.0
```

```
mkfs.acfs: volume = /dev/asm/volume1-123
```

```
mkfs.acfs: volume size = 10737418240
```

```
mkfs.acfs: Format complete.
```

See "mkfs" (Linux or UNIX) or "acfsformat" (Windows). The root privilege is not required. The ownership of the volume device file dictates who can run this command.

Oracle® Automatic Storage Management Administrator's Guide

11g Release 2 (11.2)

Question: 8

Some new non-ASM shared storage has been made available by the storage administrator, and the Oracle Grid Infrastructure administrator decides to move the voting disks, which do not reside in ASM, to this new non-ASM location. How can this be done?

- A. by running crsctl add css votedisk <path_to_new_location> followed by crsctl delete css –votedisk <path_to_old_location>
- B. by running crsctl replace css votedisk <path_to_old_location,path_to_new_location>
- C. by running srvctl replace css votedisk <path_to_old_location, path_to_new_location>
- D. by running srvctl add css votedisk <path_to_new_location> followed by srvctl delete css votedisk <path_to_old_location>

Answer: A

Explanation:

Adding, Deleting, or Migrating Voting Disks

Modifying voting disks that are stored in Oracle ASM

To migrate voting disks from Oracle ASM to an alternative storage device, specify the path to the non-Oracle

ASM storage device with which you want to replace the Oracle ASM disk group using the following command:

```
$ crsctl replace votedisk path_to_voting_disk
```

You can run this command on any node in the cluster.

To replace all voting disks not stored in Oracle ASM with voting disks managed by Oracle ASM in an Oracle

ASM disk group, run the following command:

```
$ crsctl replace votedisk +asm_disk_group
```

Modifying voting disks that are not stored on Oracle ASM:

To add one or more voting disks, run the following command, replacing the path_to_voting_disk variable with one or more space-delimited, complete paths to the voting disks you want to add:

```
$ crsctl add css votedisk path_to_voting_disk [...]
```

To replace voting disk A with voting disk B, you must add voting disk B, and then delete voting disk A.

To add a new disk and remove the existing disk, run the following command, replacing the path_to_voting_diskB variable with the fully qualified path name of voting disk B:

```
$ crsctl add css votedisk path_to_voting_diskB -purge
```

The -purge option deletes existing voting disks.

To remove a voting disk, run the following command, specifying one or more space-delimited, voting disk FUIDs or comma-delimited directory paths to the voting disks you want to remove:

```
$ crsctl delete css votedisk {FUID | path_to_voting_disk[...]}
```

Oracle® Clusterware Administration and Deployment Guide

11g Release 2 (11.2)

Question: 9

Which three statements are true about using RMAN with ASM?

- A. RMAN is the only supported method to back up database files stored in ASM.
- B. RMAN is the only supported method to back up ACFS files.
- C. RMAN can use ASM storage for backups.
- D. RMAN cannot use ASM storage for backups.
- E. Using RMAN, database files can be migrated to ASM from a file system.
- F. Using RMAN, database files cannot be moved from ASM to a file system

Answer: A, C, E

Explanation:

- RMAN is also critical to Automatic Storage Management (ASM). Since ASM is critical to GRID operations, RMAN is then critical for GRID based systems. In ASM, RMAN is responsible for tracking the ASM filenames and for performing the deletion of obsolete ASM files. ASM files cannot, however, be accessed through normal operating system interfaces; therefore, RMAN is the preferred means of copying ASM files. It is also possible to use FTP through XDB, but generally, RMAN will be

less complex. It is important to learn to use RMAN first because RMAN is the only method for performing backups of a database containing ASM files.

- Setting Initialization Parameters for ASM

The procedure for creating a duplicate database to an ASM location is similar to the procedure described in "Settings and Restrictions for OMF Initialization Parameters" on page 25-4. The difference is that you must identify the initialization parameters that control the location where files are created and set these parameters to an ASM disk group. For example, set DB_CREATE_FILE_DEST, DB_CREATE_ONLINE_DEST_n, and CONTROL_FILES to +DISK1.

- To take advantage of Automatic Storage Management (ASM) with an existing database you must migrate that database into ASM. This migration is performed using Recovery Manager (RMAN) even if you are not using RMAN for your primary backup and recovery strategy.

Oracle® Database Backup and Recovery Advanced User's Guide

Question: 10

Which two statements are true about ACFS snapshots?

- A. They can be created for ACFS file systems only if the ASM disk group hosting the ADVM volume file used by the file system has free space available.
- B. They can be created for ACFS file systems only if the ADVM volume file used by the file system has free space available.
- C. They can be created only if the ASM disk group hosting the ADVM volume used by the file system has no other ASM files contained in the disk group.
- D. They can be created when ACFS is used both on clusters and on stand-alone servers.
- E. They are accessible only on the cluster node that was used when creating the snapshot.

Answer: B, D

Explanation:

About Oracle ACFS Snapshots

Oracle ACFS snapshot storage is maintained within the file system, eliminating the management of separate storage pools for file systems and snapshots. Oracle ACFS file systems can be dynamically resized to accommodate additional file and snapshot storage requirements.

Oracle® Automatic Storage Management Administrator's Guide

11g Release 2 (11.2)

Question: 11

Your cluster is subject to a service-level agreement that allows for little scheduled down time. You want to use patching and maintenance methods that permit the Oracle Grid Infrastructure and Oracle RAC Databases to be available as much as possible. Which two techniques will work some or all of the time to provide you with minimum down time?

- A. rolling upgradeable and in place patch sets
- B. idling upgradeable and out of place patch bundles
- C. rolling upgradeable and out of place patch sets
- D. rolling upgradeable and in place patch bundles

E. rolling upgradeable and out of place one-off patches

Answer: C, D

Explanation:

- A rolling upgrade process allows one or more nodes in the cluster to have a different software version than the other nodes in the cluster for a short period of time. During the time period where cluster nodes have different software versions, the cluster is fully functional, but some administrative functions are disabled until the upgrade has completed, for example, adding new cluster resources.
- An out-of-place patch set installation with cloning has the following advantages:
Applications remain available while software is upgraded in the new ORACLE_HOME.
The configuration inside the ORACLE_HOME is retained because the cloning procedure involves physically copying the software
- In-place upgrade – You install the patch set into an existing Oracle home location. Oracle recommends that you select this option only if you do not have sufficient free disk space to perform an out-of-place upgrade, as the upgrade removes the existing Oracle installation. This patch option requires less disk space, but requires more time, and is riskier, because if you encounter an installation failure, then you must recover the entire existing Oracle home from a backup

Question: 12

Which two types of files can be stored In an ASM clustered file system?

- A. OCR and Voting Disk files
- B. data files for external tables
- C. Oracle database executable
- D. Grid Infrastructure executables
- E. data files for tablespaces
- F. archive log files

Answer: B, C

Explanation:

Oracle Automatic Storage Management Cluster File System (Oracle ACFS) is a multi-platform, scalable file system, and storage management technology that extends Oracle Automatic Storage Management (Oracle ASM) functionality to support customer files maintained outside of Oracle Database. Oracle ACFS supports many database and application files, including executables, database trace files, database alert logs, application reports, BFILEs, and configuration files. Other supported files are video, audio, text, images, engineering drawings, and other general-purpose application file data.

Notes:

- Oracle ASM is the preferred storage manager for all database files. It has been specifically designed and optimized to provide the best performance for database file types.
- Oracle ACFS is the preferred file manager for non-database files. It is optimized for general purpose files.
- Oracle ACFS does not support any file type that can be directly stored in Oracle ASM, except where explicitly noted in the documentation.

Not supported means Oracle Support Services does not take calls and development does not fix bugs associated with storing unsupported file types in Oracle ACFS.

- Starting with Oracle Automatic Storage Management 11g Release 2 (11.2.0.3), Oracle ACFS supports RMAN backups (BACKUPSET file type), archive logs (ARCHIVELOG file type), and Data Pump dumpsets (DUMPSET file type). Note that Oracle ACFS snapshots are not supported with these files.
- Oracle ACFS does not support files for the Oracle Grid Infrastructure home.
- Oracle ACFS does not support Oracle Cluster Registry (OCR) and voting files.

Question: 13

You want to reorganize the DATA diskgroup while continuing database operations. The DATA diskgroup was created using normal redundancy having one disk per failure group. The two disks used are /dev/sdal and /dev/sda2.

You plan to drop the existing disks and add the /dev/sdb1 and /dev/sdb2 disks to failure group FG_C and the /dev/sdcl and /dev/sdc2 disks to failure group FG_D.

Which procedure would you use to minimize the effect of the I/Os of this reorganization on ongoing database operations?

A.

Set rebalance power to 0 for diskgroup DATA.

Add failure group FG_C with all the /dev/sdb disks.

Add failure group FG_D with all the /dev/sdc disks.

Drop disks/dev/sda1 and /dev/sda2.

Set rebalance power to 1 for diskgroup DATA.

B.

Set rebalance power to 0 for diskgroup DATA.

Add failure group FG_C with all the /dev/sdb disks.

Add failure group FG_D with all the /dev/sdb disks.

Drop disks/dev/sda1 and /dev/sda2.

Set rebalance power to 9 for diskgroup DATA.

C.

Set rebalance power to 9 for diskgroup DATA.

Add failure group FG_C with all the /dev/sdb disks.

Add failure group FG_D with all the /dev/sdc disks.

Drop disks /dev/sda1 and /dev/sda2.

Set rebalance power to 0 for diskgroup DATA.

D.

Set rebalance power to 0 for diskgroup DATA

Drop disks /dev/sdal and /dev/sdb disks.

Add failure group FG_C with all the /dev/sdb disks.

Add failure group FG_D with all the /dev/sdc disks

Set rebalance power to 1 for diskgroup DATA.

Answer: A

Explanation:

- To control the speed and resource consumption of the rebalance operation, you can include the REBALANCE POWER clause in statements that add, drop, or resize disks.

- The ASM_POWER_LIMIT initialization parameter specifies the default power for disk rebalancing in a disk group. The range of values is 0 to 1024. The default value is 1. A value of 0 disables rebalancing. Higher numeric values enable the rebalancing operation to complete more quickly, but might result in higher I/O overhead and more rebalancing processes.
- Failure groups are used to place mirrored copies of data so that each copy is on a disk in a different failure group. The simultaneous failure of all disks in a failure group does not result in data loss.
- You define the failure groups for a disk group when you create an Oracle ASM disk group. After a disk group is created, you cannot alter the redundancy level of the disk group. If you omit the failure group specification, then Oracle ASM automatically places each disk into its own failure group, except for disk groups containing disks on Oracle Exadata cells. Normal redundancy disk groups require at least two failure groups. High redundancy disk groups require at least three failure groups. Disk groups with external redundancy do not use failure groups.

Oracle® Automatic Storage Management

Thank You for trying 1Z0-058 PDF Demo

To Buy Latest 1Z0-058 Full Version Download visit link below

<https://www.certkillers.net/Exam/1Z0-058>

Start Your 1Z0-058 Preparation

[Limited Time Offer] Use Coupon “CKNET” for Further discount on your purchase. Test your 1Z0-058 preparation with actual exam questions.

<https://www.certkillers.net>